

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND SYSTEM FOR MANAGING
INVERSE MULTIPLEXING OVER ATM**

First Named Inventor:

Cetin Seren

Prepared by:
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, CA 90025-1026
(408) 720-8300

"Express Mail" mailing label number EL627469505US

Date of Deposit December 28, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Geneva Walls

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

METHOD AND SYSTEM FOR MANAGING INVERSE MULTIPLEXING OVER ATM

Field of the Invention

The present invention pertains to communications and networking. More particularly, the invention relates to the usage of resources in networking devices.

5 Background of the Invention

Packet switching communication networks are commonly employed to transfer digital information over long distances. One type of packet switching communication network is an asynchronous transfer mode communication network. An asynchronous transfer mode communication network enables the transfer of digital voice information,
10 digital video information and digital data over short or long distances via common carrier communication links. Such a communication network enables a wide variety of communication devices to share common carrier communication links on a demand driven basis. For example, such asynchronous transfer mode communication networks are commonly employed for a wide area network (WAN) communication.

15 The common carrier communication links employed in such asynchronous transfer mode communication networks typically include relatively low speed metal wire based communication links. One such type of relatively low speed common carrier communication link available in North America is the T1 communication link. A typical T1 communication link provides a maximum data rate of 1.544 megabits per second.
20 Similarly, the E1 common carrier communication links available in Europe provide a data rate of 2.048 megabits per second.

In addition, the common carrier communication links employed in such asynchronous transfer mode communication networks usually include higher speed communication links. An example of such a high-speed common carrier communication link available in North America is the T3 communication link, which provides a maximum data rate of 45 megabits per second. An example of a high-speed common carrier communication link available in Europe is the E3 communication link, which provides a maximum data rate of 34 megabits per second.

Prior asynchronous transfer mode communication networks that require high bandwidth communication links typically employ T3 or E3 communication links rather than the lower speed T1 or E1 communication links. Unfortunately, the subscription rates for such T3 common carrier communication links are usually much higher than the subscription rates for T1 or E1 links. Moreover, in North America such T3 common carrier communication links are not as widely available as T1 common carrier communication links. Worse yet, in Europe E3 common carrier communication links are rarely available and E1 common carrier communication links have only recently become widely available.

As a consequence, communication services requiring such high bandwidth communication links typically require a high premium payment for such high speed common carrier communication links. Moreover, such required high-speed communication links may not be available in some areas for such an asynchronous transfer mode communication network.

In addition, some prior asynchronous transfer mode communication networks require communication links with a bandwidth higher than is provided by low speed

common carrier communication links such as T1 or E1 but do not require the high bandwidth of the T3 or E3 common carrier communication links. Unfortunately, communication services requiring such medium level bandwidth communication links must typically employ expensive high-speed common carrier communication links even though the full bandwidth range is not needed. For example, an asynchronous transfer mode communication network requiring a maximum data rate of 10 or 20 megabits per second must usually employ high speed T3 or E3 common carrier communication links. Such high-speed common carrier links unnecessarily increase the cost of communication for subscribers to such an asynchronous transfer mode communication networks.

In an Inverse Multiplexing for ATM (IMA) system, ATM cell traffic is transported by means of time-division multiplexing over several channels (typically T1 or E1 data links). In a cell based IMA system, these ATM cells are sent on each channel in a round-robin fashion as depicted in Figures 1A-1C.

The receiving IMA device must reconstruct its output stream from cells received over the constituent channels, in such a way that cell sequence integrity is preserved.

Referring to the figures, the basic function of IMA device is to work in pairs to take an ATM cell stream coming from the ATM layer, send it over the multiple links by spreading cells over the available links and ensure that the initial cell stream can be retrieved at the far end. Thus the IMA preferably makes the ATM traffic transparent to the ATM layer over multiple links in use. As far as the ATM layer is concerned, it should only see a pipe (can be considered as a virtual link) whose rate is now the sum of the multiple link rates. It is assumed that each link is run in clear-mode without the presence

SUMMARY OF THE INVENTION

What is disclosed is a method and system for managing inverse multiplexing over ATM. In one embodiment, a method for use in a digital communications network, comprises monitoring a plurality of links to determine state changes of the links. An
5 IMA-ID check is enforced when an insufficient links state is reached. The IMA-ID check is relaxed when all the links are in an error state. An IMA-ID check is re-enforced when at least one link of the plurality of links recovers from the error state.

Other features and advantages of the present invention will be apparent from the accompanying drawings and from the detailed description that follows below.

10

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitations in the figures of the accompanying drawings, in which like references indicate similar elements, and in which:

5 **Figure 1A** shows a prior art IMA system;

Figure 1B shows a prior art IMA system with one link in loopback;

Figure 1C shows a prior art IMA system with one link in loopback and the remaining links broken;

Figure 2 shows an enhanced IMA system;

10 **Figure 3** shows a block diagram of a group restarter;

Figure 4 shows a flow diagram of process performed by group restarter to selectively allow the restart of IMA groups; and

Figure 5 shows a process performed by a group restarter for making a group ID persistent.

DETAILED DESCRIPTION

A method and system for managing IMA. As described in detail below, in one embodiment of the present invention, a method for use in a digital communications network, comprises monitoring a plurality of links to determine state changes of the links. An IMA-ID check is enforced when an insufficient links state is reached. The IMA-ID check is relaxed when all the links are in an error state. An IMA-ID check is re-enforced when at least one link of the plurality of links recovers from an error state .

In one embodiment, the present method and system address when far end IMA devices may change, or when office wires may be misconnected to a far end device. As described in detail below, these unintended data transmissions are corrected. In another embodiment, the present method and system address what occurs when links are put into a loopback state where the IMA group does not know if it should be communicating with itself or if it should be communicating with a far end IMA device. As described in detail below, the data transfer between a near end device in loopback and an intended far end IMA device are assured.

Although embodiments of the present invention are described as having both software and hardware elements, alternative embodiments may be all hardware, all software, or a combination of each. The software implementing the present invention can be stored in RAM, a mass storage device available through disk interface, or other storage medium accessible to a CPU. This software may also be resident on an article of manufacture comprising a computer usable mass storage medium or propagated digital signal having computer readable program code embodied therein and being readable by the mass storage medium.

Figure 1A shows a prior art fully operational IMA system 100, such as that described in *Inverse Multiplexing for ATM (IMA) Specification*, ATM Forum Technical Committee, December 1998, herein incorporated by reference. IMA Group 101 has four links (IMA lines 121-124) to IMA group 111. Lines 121-124 may be high-speed data lines such as T1, E1, J1, T3, E3, or J3 connections. Lines 121-124 are bi-directional. In one example, data cells are transmitted from IMA group 101 in a round-robin multiplexed fashion over lines 121-124. IMA group 111 reconstructs the multiplexed cells. IMA groups 101, 111 may be IMA entities, such as, routers or concentrators for pumping data into a switch.

Figure 1B shows a prior art IMA system 130 with one link in loopback. Line 151 may be a faulty line and placed in loopback. When line 151 is in loopback, cells from IMA Group 131 are not transmitted over line 151. Typically a system operator places a line into loopback.

Figure 1C shows a prior art IMA system 150 with one link 171 in loopback and the remaining links 171-174 broken. For example, lines 171-174 may have been physically cut or damaged. IMA groups 151 and 161 go into a "failed" state and alarms occur. When IMA Groups 151, 161 restart, line 171 in loopback comes out of alarm faster than lines 172-174. IMA group 161 is completely cut off, and IMA group 151 "talks" to itself via loopbacked link 171. Since line 171 will be the first line to become operational after restart, an IMA connection is established with itself. However, lines 172-174 will not have matching IMA parameters and be declared unusable. Thus, upon restart, instead of synchronizing with IMA Group 161, IMA group 151 synchronizes with itself. However, loopbacks are an essential part of maintenance.

Figure 2 shows an enhanced IMA system 200. System 200 includes a Near End IMA Group 205 with links 201 to 204. Far End IMA Group 209 is also connected to links 201 to 204. IMA Groups 205, 209 may be IMA entities, such as routers or concentrators for pumping data into a switch. Links 201-204 may be high-speed data lines, such as E1, T1, J1, E3, T3, or J3 data lines. IMA Group 205 includes Group Restarter 206. Group restarter 206 facilitates IMA Group start-ups as will be discussed in detail below. Also included in IMA group 205, is a group state machine (GSM) 207 which monitors if the links enter an insufficient links state. Although each IMA group 205, 209 is depicted as having a Group Restarter 206, in alternate embodiments only one Group Restarter 206, 208 may be implemented according to the present method and system.

Figure 3 shows a block diagram of a group restarter 206. One embodiment of group restarter 206 comprises a system bus 220 for communicating information, and a processor 210 coupled to bus 220 for processing information. Group restarter 206 further comprises a random access memory (RAM) or other dynamic storage device 225 (referred to herein as main memory), coupled to bus 220 for storing information and instructions to be executed by processor 210. Main memory 225 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 210. Group restarter 206 also may include a read only memory (ROM) and/or other static storage device 226 coupled to bus 220 for storing static information and instructions used by processor 210.

A data storage device 227 such as a magnetic disk or optical disc and its corresponding drive may also be coupled to group restarter 206 for storing information

and instructions. Group restarter 206 can also be coupled to a second I/O bus 250 via an I/O interface 230. A plurality of I/O devices may be coupled to I/O bus 250, including a display device 243, an input device (e.g., an alphanumeric input device 242 and/or a cursor control device 241). For example, video news clips and related information may be presented to the user on the display device 243.

The communication device 240 is for accessing other computers (servers or clients) via a network. The communication device 240 may comprise a modem, a network interface card, or other well known interface device, such as those used for coupling to Ethernet, token ring, or other types of networks.

Figure 4 shows a flow diagram of process performed by group restarter 206 to selectively allow the restart of IMA groups. The process of **Figure 4** may establish new connections with an IMA group when it appears that a new IMA group is connected. Furthermore, the process of **Figure 4** will maintain the links between two IMA groups when a link may be in a loopback state. Prior to beginning the process of **Figure 4**, a group state machine (GSM) 207 is started and initialized. The IMA ID check is enforced for all links and the GSM 207 enters an insufficient link state. This may occur, for example when a new device is about to be connected to a network as a far end IMA group. Then, the individual links are monitored by a link state machine (LSM) as described below.

The process starts at block 400. At processing block 410, group restarter 206 monitors operational links 201-204. At decision block 420, links 201-204 are monitored to detect if all the links are in an error state. If not all links are in error, then flow continues to decision block 430. If all the links in group 205 are in error, then flow

continues to decision block 440. At decision block 440, group restarter 206 checks to see if an IMA-ID check has already been enforced. If the check has not been enforced already, then flow passes to start block 400. If the check has been enforced, then flow passes to processing block 460. By enforcing an IMA-ID check, a link is only accepted
5 if the far end IMA-ID matches a predetermined (expected) IMA-ID derived from the same link as described in detail below. At processing block 460, the IMA-ID check is relaxed and flow returns to start block 400. By relaxing an IMA-ID check, the IMA-ID check requirement is disregarded for a particular link. The IMA-ID check is relaxed to minimize the alarms that occurs when all links are in error. Thus, block 460 relaxes an
10 IMA-ID check in certain cases, for example, when an IMA Group is replaced. After relaxing the IMA-ID check, flow continues back to start block 400.

If all links are not in error and flow continued to decision block 430, links 201-204 are monitored to detect if any one recovers from an error. If all links are operational and there is no error recovery, then flow passes back to start block 400. If group restarter
15 206 detects that a link 201-204 has recovered from an error, flow continues to processing block 450. At processing block 450 in the IMA-ID check is enforced as follows. If the IMA-ID does not match, the link is broken and made unusable. If the IMA-ID matches, the link is maintained. Then, flow passes back to start block 400. Thus, an IMA-ID check is enforced in cases where a link 201-204 recovers from an
20 error and is still connected to the same IMA group 206. No alarms occur.

Figure 4 may be modified to accommodate special conditions, such as when lines are in loopback. To detect if a line is in loopback, an additional condition is added to decision blocks 410, which is to check that the IMA group IDs are different.

Every IMA Group has an identifier (ID). Cells are transmitted between two IMA groups. One out of every 128 cells transmitted between two IMA groups includes an IMA Control Protocol Cell (ICP). The ICP is a 53 byte cell that contains the IMA Group's ID (IMA-ID). The IMA-ID may not be unique since it is established independently of neighboring IMA Groups. Furthermore, IMA-ID's are only 8 bits in length. Every 128 cells, the IMA-ID is checked to see if it matches an expected IMA-ID as established during initialization of the links. The IMA-ID is checked by a Frame Synchronization Mechanism (FSM).

The process of **Figure 4** institutes two new actions, a relax IMA-ID check action and an enforce IMD-ID check action. These two additional actions will allow IMA groups to be resynchronized on a per group basis. IMA groups that are resynchronized establish data connectivity with a far end IMA device. In addition, an IMA Group ID may be made persistent. For example, an IMA Group's ID will be expected to be the same, both before and after resetting the line card upon which the IMA Group resides.

Figure 5 shows a process performed by a group restarter for making a group ID persistent. Group restarter 206 learns the IMA group ID of IMA group 209, by reading it from an ICP cell as described above at processing block 510. At processing block 520, the IMA group ID is made persistent by storing in storage 227 the ID value obtained at block 510 at block. At block 530, only links are used that have an IMA group with the same ID stored in storage 227 will be used. At processing block 540, non-matching links are placed in an unusable state. Flow continues back to processing block 520 as described above. The process of **Figure 5** is used when an IMA group has not been replaced.

In another embodiment, a process may be performed by a group restarter for relearning a group ID. The IMA group 205 is restarted. Then, the IMA-group ID of group 209 is learned from the ICP cell as described above. The IMA group ID is made persistent by storing in storage 227 the ID value learned. The new IMA ID is provided to the FSM for the IMA-ID check. All operational links are then established.

The IMA groups were designed so that once a link is established it should never change. That rigidity was invoked to divert security breaches. However, when all links are down, in the case where the links have been reconnected to another IMA entity, a new far-end IMA device should be accepted. The process does not restart the whole group or reset the card on which the IMA entity resides, because there will be a 5-10 minute system outage resulting in data loss.

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawing are, accordingly, to be regarded in an illustrative rather than a restrictive manner.